

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Richard Wiss et al.	Examiner: Pham, Khanh B
Serial No.: 10/605,154	Art Unit: 2166
Filed: September 11, 2003	APPEAL BRIEF (Amended)
For: Database System Providing Improved Methods For Data Replication	

Mail Stop Appeal
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

BRIEF ON BEHALF OF RICHARD WISS ET AL.

This is an appeal from the Final Rejection mailed August 23, 2006 in which currently-pending claims 1-43 stand finally rejected. Appellant filed a Notice of Appeal on December 28, 2006 (as indicated by return of a confirmation postcard marked "OIPE DEC 28 2006"). This brief is submitted electronically in support of Appellant's appeal.

TABLE OF CONTENTS

1.	REAL PARTY IN INTEREST	3
2.	RELATED APPEALS AND INTERFERENCES	3
3.	STATUS OF CLAIMS.....	3
4.	STATUS OF AMENDMENTS.....	3
5.	SUMMARY OF CLAIMED SUBJECT MATTER.....	4
6.	GROUND OF REJECTION TO BE REVIEWED	7
7.	ARGUMENT	7
	A. First Ground: Claims 1-8, 10-23, 25-32, and 34-43 rejected under Section 102	7
	B. Second Ground: Claims 9, 24 and 33 rejected under Section 103	13
8.	CONCLUSION	14
9.	CLAIMS APPENDIX	16
10.	EVIDENCE APPENDIX	23
11.	RELATED PROCEEDINGS APPENDIX.....	24

1. REAL PARTY IN INTEREST

The real party in interest is assignee Sybase, Inc., located at One Sybase Drive, Dublin, CA 94568.

2. RELATED APPEALS AND INTERFERENCES

There are no appeals or interferences known to Appellant, the Appellant's legal representative, or assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

3. STATUS OF CLAIMS

The status of all claims in the proceeding is as follows:

Rejected: Claims 1-43

Allowed or Confirmed: None

Withdrawn: None

Objected to: None

Canceled: None

Identification of claims that are being appealed: Claims 1-43

An appendix setting forth the claims involved in the appeal is included as the last section of this brief.

4. STATUS OF AMENDMENTS

One Amendment has been filed in this case. Appellant filed an Amendment on May 3, 2006, in response to a non-final Office Action dated January 3, 2006. In the Amendment, the pending claims were amended in a manner which Appellant believes clearly distinguished the claimed invention over the art of record, for overcoming the art rejections. In response to the Examiner's Final Rejection dated August 23, 2006, Appellant filed a Notice of Appeal. Appellant has chosen to forgo filing an Amendment After Final which might further limit Appellant's claims, as it is believed that further amendments to the claims are not warranted in view of the art. Accordingly, no Amendments have been entered in this case after the date of the Final Rejection.

5. SUMMARY OF CLAIMED SUBJECT MATTER

As to Appellant's **First Ground** for appeal, Appellant asserts that the art rejection under **Section 102(e)** relying on **Shih** fails to teach or suggest all of the claim limitations of Appellant's rejected claims 1-8, 10-23, 25-32, and 34-43, where the claimed invention is set forth in the embodiment in **independent claim 1**: a method for replicating a transaction from a primary database to a replicate database while the replicate database remains available for use (see, e.g., Applicant's Specification at **Paragraph [0056]** and **Fig. 3, at 310, 330**), the method comprising: recording information about a transaction being performed at a primary database in a transaction log (see, e.g., Applicant's Specification at **Paragraphs [0069] and [0070]** and **Fig. 4, at 403**); synchronously copying the information about the transaction in the transaction log to a mirrored transaction log, so as to create at the replicate database an exact copy of the transaction log (see, e.g., Applicant's Specification at **Paragraph [0071]** and **Fig. 4, at 404**); generating a reconstructed transaction based on the information about the transaction copied to the mirrored transaction log (see, e.g., Applicant's Specification at **Paragraph [0072]** and **Fig. 4, at 405, 406**); and applying the reconstructed transaction at the replicate database while the replicate database remains available for use (see, e.g., Applicant's Specification at **Paragraphs [0073] - [0075]** and **Fig. 4, at 407-409**).

Appellant further asserts that the art rejection relying on **Shih** fails to teach or suggest all of the claim limitations of Appellant's claimed invention, where the claimed invention comprises the embodiment set forth in **independent claim 18**: a system for replicating transactions (see, e.g., Applicant's Specification at **Paragraph [0056]** and **Fig. 3, at 300, 310, 330**) from a source database to a standby database that comprises: a source database having a transaction log, the transaction log for recording log records for transactions performed at the source database (see, e.g., Applicant's Specification at **Paragraphs [0069] and [0070]** and **Fig. 4, at 403**); a mirrored transaction log for recording mirror copies of the log records for transactions performed at the source database, so as to create at the standby database an exact copy of the transaction log (see, e.g., Applicant's Specification at **Paragraph [0071]** and **Fig. 4, at 404**); a file mirroring module for synchronously replicating log records from the transaction log to the mirrored transaction log as transactions are performed at the source database (see, e.g., Applicant's

Specification at **Paragraph [0072]** and **Fig. 4, at 405**); a log reader module for reading log records in the mirrored transaction log and reconstructing transactions for application at the standby database based upon log records in the mirrored transaction log (see, e.g., Applicant's Specification at **Paragraph [0072]** and **Fig. 4, at 406**); and a distribution module for applying the transactions reconstructed by the log reader module at the standby database (see, e.g., Applicant's Specification at **Paragraphs [0073] - [0075]** and **Fig. 4, at 407-409**).

Appellant further asserts that the art rejection relying on **Shih** fails to teach or suggest all of the claim limitations of Appellant's claimed invention, where the claimed invention comprises the embodiment set forth in **independent claim 30**: a method for replicating a database operation from a first database to a second database while making the second database available for decision support purposes (see, e.g., Applicant's Specification at **Paragraph [0056]** and **Fig. 3, at 310, 330**), the method comprises steps of: as a database operation is performed at the first database, generating at least one log record characterizing the operation (see, e.g., Applicant's Specification at **Paragraphs [0069]** and **[0070]** and **Fig. 4, at 403**); synchronously recording the at least one log record in a first log associated with the first database and a second log associated with the first log, so that the second log comprises an exact copy of the first log (see, e.g., Applicant's Specification at **Paragraph [0071]** and **Fig. 4, at 404**); and while the second database is available for decision support purposes, replicating the operation performed at the first database at the second database by performing the substeps of: constructing a replicate operation based, at least in part, on the at least one log record in the second log; and applying the replicate operation at the second database (see, e.g., Applicant's Specification at **Paragraph [0072]** and **Fig. 4, at 405, 406**; and **Paragraphs [0073] - [0075]** and **Fig. 4, at 407-409**).

Appellant further asserts that the art rejection relying on **Shih** fails to teach or suggest all of the claim limitations of Appellant's claimed invention, where the claimed invention comprises the embodiment set forth in **independent claim 43**: a method of the present invention for replicating transactions from a primary database to a replicate database while the replicate database remains available for use (see, e.g., Applicant's Specification at **Paragraph [0056]** and **Fig. 3, at 310, 330**), the method comprises steps

of: recording log records for transactions being performed at a primary database in a primary transaction log (see, e.g., Applicant's Specification at **Paragraphs [0069] and [0070]** and **Fig. 4, at 403**); creating a mirrored transaction log, the mirrored transaction log comprising an exact copy of the log records in the primary transaction log (see, e.g., Applicant's Specification at **Paragraph [0071]** and **Fig. 4, at 404**); generating reconstructed transactions based on the copies of the log records in the mirrored transaction log (see, e.g., Applicant's Specification at **Paragraph [0072]** and **Fig. 4, at 405, 406**); and applying the reconstructed transactions at the replicate database while the replicate database remains available for use (see, e.g., Applicant's Specification at **Paragraphs [0073] - [0075]** and **Fig. 4, at 407-409**).

As to Appellant's **Second Ground** for appeal, Appellant asserts that the art rejection under **Section 102(e)** relying on **Shih** fails to teach or suggest all of the claim limitations of Appellant's rejected claims 9, 24 and 33. Regarding **dependent claim 9**, the claimed invention is set forth in the claimed invention in the embodiment in parent and **independent claim 1** (the mapping of which is shown above under Appellant's First Ground for appeal, and which hereby incorporated by reference). Regarding **dependent claim 24**, the claimed invention is set forth in the claimed invention in the embodiment in parent and **independent claim 18** (the mapping of which is shown above under Appellant's First Ground for appeal, and which hereby incorporated by reference). Regarding **dependent claim 33**, the claimed invention is set forth in the claimed invention in the embodiment in parent and **independent claim 30** (the mapping of which is shown above under Appellant's First Ground for appeal, and which hereby incorporated by reference). For Appellant's argument under the Second Ground for appeal, Appellant **additionally argues** based on **dependent claims 9, 24, and 33** which include additional features of file block level limitations, pertaining to synchronously copying by replicating at a file block level the information about the transaction in the transaction log to the mirrored transaction log (see, e.g., Applicant's Specification at **Paragraphs [0071], [0089] - [0092], and [0100] - [0118]** and **Figs. 6A-B, at 601 - 610**).

6. GROUNDS OF REJECTION TO BE REVIEWED

The grounds of rejection to be reviewed on appeal are:

(1st) Whether claims 1-8, 10-23, 25-32, and 34-43 are unpatentable under 35 U.S.C. 102(e) as being anticipated by Shih et al. (US 6,615,223 81); and

(2nd) Whether claims 9, 24 and 33 are unpatentable under 35 U.S.C. 103(a) as being obvious over Shih as applied above, and in view of Riedel et al.

7. ARGUMENT

A. First Ground: Claims 1-8, 10-23, 25-32, and 34-43 rejected under Section 102

1. General

Under Section 102, a claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in the single prior art reference. (See, e.g., MPEP Section 2131.) As will be shown below, the reference fails to teach each and every element set forth in claim 1, as well as other claims, and therefore fails to establish anticipation of the claimed invention under Section 102.

2. Claims 1-8, 10-23, 25-32, and 34-43

Claims 1-8, 10-23, 25-32, and 34-43 are rejected under 35 U.S.C. 102(e) as being anticipated by Shih et al. (US 6,615,223 81), hereinafter "Shih". The Examiner's rejection of claims 1, 16, and 17 is representative:

As per claims 1, 16-17, Shih teaches a method and computer readable medium for replicating a transaction from a primary database to a replicate database while the replicate database remains available for use (Col. 9 lines 15-42), the method comprising:

"recording information about a transaction being performed at a primary database in a transaction log" at Col. 9 lines 15-25 and Fig. 3;

"synchronously copying the info about the transaction in the transaction log to a mirrored transaction log" at Col. 9 lines 28-60;

"generating a reconstructed transaction based on the information about the transaction copied to the mirrored transaction log"

at Col. 10 lines 10-45;

"applying the reconstructed transaction at the replicate database while the replicate database remains available for use" at Col. 9 lines 28-42.

As will be shown below, Appellant's invention may be distinguished on a variety of grounds.

Shih does describe a sync replication system, and as a result shares some features in common with Appellant's system. However, Shih's approach has distinct disadvantages that Appellant's invention overcomes. (Appellant in fact had considered such a solution and discarded it due to lack of scalability and performance.) The discussion which follows identifies deficiencies of Shih and highlights features of Appellant's invention that overcome those deficiencies, as well as highlights where such features are recited as specific claim limitations present in Appellant's amended claims.

As a core architectural difference, the Shih system does **not** have an **exact copy** (mirror image) of the log (Appellant's claim limitation of "mirrored transaction log") at replicate sites, as log archiving and truncating operate independently of one another on the two systems. Consider the following teaching from Shih:

When a change request 12 is received at first replication site 2, server 10 issues change instruction 16 to implement the change request 12. **The change instruction 16 takes into account the exact schema organization of the data object to be changed. Thus, the change instruction is schema-specific, and in a heterogeneous environment cannot simply be sent to all remote replication sites to replicate the data change, since the schema and/or system configuration of the remote replication sites may be entirely different** than the schema and system configuration of local replication site 2.

According to the invention, server 10 translates either change instruction 16 or change request 12 into a schema and system independent change record 20. Change record is in a generic format that is consistent and recognizable across all replication sites in the system. In the normal contemplated usage of the invention, change record 20 comprises change information that is focussed upon the specific data to be added, deleted, or modified by the change request 12, and does not contain information

regarding the schema organization of the data at the originating replication site.

(Shih, at col. 5, lines 5-25, emphasis added.)

In order for Shih's system to support schema independence between primary and replicates, it performs the above quoted translation to write schema independent log records. (As an aside, commercial database systems generally do not write schema independent log records, thereby limiting customer appeal for Shih's approach.) In fact recall that Appellant's described system is one that operates in real-time (e.g., for supporting OLTP services), and therefore the additional processing required for writing schema independent log records would impact primary database performance in a manner that would be entirely unacceptable for such time-critical applications.

These differences between Appellant's claimed approach and that of Shih are not merely theoretical differences but ones that have important practical implications. Consider the following teaching from Appellant's specification:

In operation, the file mirroring module 315 replicates the log records from the primary log 313 on the primary server 310 to the mirrored log 332 which is typically located at the same site as the standby server 330. The file mirroring module 315 is a disk mirroring technology that replicates the log record data very quickly and efficiently. In addition, in the currently preferred embodiment of the present invention the file mirroring module operates synchronously. **When log records are written from the primary database 311 to the primary log 313, copies of these log records are written at the same time to the mirrored primary log 332 at the standby location. This guards against any loss of data in the event of a problem with the primary database server.**

(Appellant's specification, Paragraph [0060], emphasis added.)

Importantly, the design choice of Shih introduces a major disadvantage: if sync write fails, the Shih system cannot recover and must either refuse updates at the primary or completely re-materialize replicates.

Appellant's system, in contrast, does not perform Shih's resource-intensive approach of translating records. Appellant's claimed approach is one that preserves the

original log file format and, thus, can work with any database system without changing the underlying database log writing function. This approach has a very important advantage. In the event that the primary database stops working, logical replication continues to be applied to the replicate database based on transactions in the mirrored primary database transaction log, enabling all of the database operations applied to the primary database to also be applied to the replicate database. This guarantees that no database operations are lost in the event of loss or damage to the primary database.

Appellant's independent claims was previously amended to highlight these distinctions. For example, independent claim 1 was amended to include the following claim limitation (shown in amended form):

synchronously copying the information about the transaction in the transaction log to a mirrored transaction log, so as to create at the replicate database an exact copy of the transaction log;

(All of Appellant's other independent claims were amended in an analogous manner.) As shown, the claim limitation requires the creation of an exact copy or mirror image of the transaction log at the replicate database. The Shih system, as described, does not teach or suggest this approach, and in fact at best Shih's translation approach appears to teach away from Appellant's claimed approach.

In the second (final) Office Action, the Examiner simply responded by dismissing the foregoing distinctions out of hand. The Examiner appears to not understand or fully appreciate that in disk mirroring a mirrored copy, such as the "mirrored transaction log" in claim 1, is an EXACT replicate backup copy for archival purposes. Significantly, mirrored copies are created using block-level operations, completely without concern as to what the data means within any given block. Using disk mirroring technique, therefore, one may create a replicate that is a byte-for-byte, block-for-block exact (physically identical) copy of the primary. A primary database, on the other hand, is a live copy created and maintained by a high-level software-level process (e.g., database server software) in a manner such that the primary is constantly changing (e.g., as a result of database transaction-based operations); the primary itself is not suitable as an archival

copy due to its constantly changing nature (i.e., the software-level process controlling the primary is constantly changing the data). The Examiner treats these two as being the same or fully interchangeable, but they are not.

Once a primary database has been mirrored to a replicate site (i.e., mirrored copy), any software process accessing the mirrored copy (e.g., for DSS-style reporting) has no control over the primary database. At the same time, the primary database continues to constantly change as a result of transactions posted to it (e.g., by database server software). Significantly, if a process were to attempt a block-by-block read of the primary's data (i.e., read the primary in a sequential manner), there is no guarantee the process would correctly read the data. For example, while a process was sequentially reading what it thought was a correct block sequence of the primary, the database server software may invalidate or rewrite one of the blocks such that the sequential read simply yields inconsistent or corrupt data. Therefore, in a system using file-based mirroring technique, one cannot even use the database at the replicate site. Instead, the copy resides at the replicate site as a mirrored backup for disaster recovery. Should a disaster occur at the primary, the mirrored backup copy can be used to restart the system (which may then continue to function in a normal manner).

The impetus for Appellant's invention is to create a system where one could use disk mirroring (because of its benefits), but do so in a manner that would allow one to bring the replicate (mirrored) database online. In order to appreciate Appellant's invention, one should understand why disk mirroring is even used in the first place. Disk mirroring is a block-level technique that provides a very fast solution to the problem of synchronously writing to remote files. A solution that attempts to do this at a software (i.e., higher) level (e.g., using distributed database technique), on the other hand, would operate fairly slowly because such a system would be required to do in essence a two-phase commit (i.e., to know that the data had been written at both the primary and replicated sites). If one were to attempt such a software-based approach over long distances, one would face a multitude of issues, including limitations imposed by the speed of light. Accordingly, such a software-based approach is impractical as a disaster recovery technique for a production database.

The Examiner underestimates the issues and problems (documented in Appellant's

prior remarks) that are encountered when trying to read a mirrored copy and manage the truncation of both primary and mirrored copies. The Examiner supposes that disk mirroring could easily be incorporated into Shih's system, but that is an incorrect assumption. Whatever process controls the primary database owns it. Therefore if a second process comes along and attempts to read the primary, it cannot simply proceed with a sequential read of file blocks because the first process will in fact be changing the data out from under it. For example, consider when a database server at the primary site truncates a transaction log file. Once the log is truncated, subsequent blocks (i.e., beyond the point of truncation) are freed and reuse for storing other data. Hence, a second process (i.e., other than the database server) which is attempting a sequential read of the transaction log file has no control over that log file and no control over how the first process (e.g., database server) writes data to that log file.

In Shih's diagram pasted by the Examiner into the final Office Action, garbage collection is done independently on the both sides. As a result, there is absolutely no way that that can be done using disk mirroring technique. The reason is as follows. At Shih's replicate site when garbage collection is occurring, the replicate site is actually writing to the data. In a disk mirroring-based system, in contrast, the replicate site never writes to the data in any sort of independent manner. Instead, it simply receives replicated blocks of data from the primary. All writing to the data is (absolutely) controlled by the primary. In Shih's case, the same information is being written to two separate files however -- importantly -- the files are not mirror (identical) copies, such as required by Appellant's claimed invention. This fact is made explicitly clear in Shih, as the two sites are independently performing garbage collection. In disk mirroring, since all writing to the data is controlled by the primary, the way to perform "garbage collection" is to go back to the primary site and cause it (not the replicate) to truncate the log file. It is that truncation that gets mirrored to the replicate site. Again, all control of the data is from the primary site.

For the reasons stated, Shih's system cannot reproduce Appellant's claimed invention simply by combining his system with mirror replication. In view of the foregoing remarks and clarifying prior amendments made to the claims, it is believed that the claims distinguish over Shih. Accordingly, it is respectfully requested that the

Examiner's rejection under Section 102(e) not be sustained.

B. Second Ground: Claims 9, 24 and 33 rejected under Section 103

1. General

Under Section 103(a), a patent may not be obtained if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the subject matter pertains. To establish a prima facie case of obviousness under this section, the Examiner must establish: (1) that there is some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings, (2) that there is a reasonable expectation of success, and (3) that the prior art reference (or references when combined) must teach or suggest all the claim limitations. (See e.g., MPEP 2142). The references cited by the Examiner fail to meet these conditions.

2. Claims 9, 24 and 33

Claims 9, 24 and 33 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Shih as applied to claims 1-8, 10-23, 25-32, 34-43 above, and in view of Riedel et al. ("When Local Becomes Global: An Application Study of Data Consistency in a Network World"), hereinafter "Riedel." Here, the Examiner essentially repeats the rejection above (under Shih) but adds Riedel for the teaching of "said synchronously copying step includes replicating at a file block level."

The claims are believed to be allowable for at least the reasons stated above pertaining to Shih. To the point, Shih does not teach or suggest Appellant's claim limitations that require the creation of an exact copy or mirror image of the transaction log at the replicate database. The Examiner has cited nothing in Riedel that remedies that deficiency.

Further, the claims are believed to be allowable for the following additional reasons. As noted above, to establish a prima facie case of obviousness under Section

103, the Examiner must establish (among other things) that there is some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings, and that there is a reasonable expectation of success. (See e.g., MPEP 2142). Appellant has invented an approach to read the mirrored log records in such a way as to guarantee consistency. This is done by controlling the primary database in such a way as to keep it from altering file blocks that have not been processed and by detecting when data is inconsistent. Internally (e.g., in the preferred embodiment), this requires a control file block loop to control the log archiving/log truncation at the primary database, which is of course absent from Shih's system. Merely strapping file block replication onto Shih's system does not reproduce Appellant's claimed approach. Importantly, the combined references provide no teaching or suggestion that would support the use of file block replication technique in Shih's system.

Shih in fact appears to instead be relying on Operating System file I/O and thus teaches away from the combination suggested by the Examiner. Furthermore, the Shih system would need to be substantially modify to incorporate the file block replication combination with Riedel (e.g., tending to such details as controlling the log archiving/log truncation at the primary database, as described for Appellant's system above), in order to create a system that could work in a manner (remotely) related to Appellant's system. Neither Shih nor Riedel provides any teaching, suggestion, or other motivation in that regard, and at best one would need to borrow heavily from the teachings of Appellant's specification in order have any reasonable expectation of success.

In view of the foregoing remarks, it is respectfully submitted that the claims distinguish over the combination of Shih and Riedel (especially in view of prior amendments to Appellant's base independent claims). Accordingly, it is respectfully requested that the Examiner's rejection under Section 103 not be sustained.

8. CONCLUSION

The present invention greatly improves the ease and efficiency of the task of managing mirror copies, in support of data replication. It is respectfully submitted that

the present invention, as set forth in the pending claims, sets forth a patentable advance over the art.

In view of the above, it is respectfully submitted that the Examiner's rejections under 35 U.S.C. Section 102 and 103 should not be sustained. If needed, Appellant's undersigned attorney can be reached at 408 884 1507. For the fee due for this Appeal Brief, please refer to the attached Fee Transmittal Sheet. This Brief is submitted electronically.

Respectfully submitted,

Date: July 26, 2007

/John A. Smart/

John A. Smart; Reg. No. 34,929
Attorney of Record

408 884 1507
815 572 8299 FAX

9. CLAIMS APPENDIX

1. (Previously presented) A method for replicating a transaction from a primary database to a replicate database while the replicate database remains available for use, the method comprising:

recording information about a transaction being performed at a primary database in a transaction log;

synchronously copying the information about the transaction in the transaction log to a mirrored transaction log, so as to create at the replicate database an exact copy of the transaction log;

generating a reconstructed transaction based on the information about the transaction copied to the mirrored transaction log; and

applying the reconstructed transaction at the replicate database while the replicate database remains available for use.

2. (Original) The method of claim 1, wherein said transaction includes a selected one of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.

3. (Original) The method of claim 1, wherein said recording step includes recording at least one log record about the transaction in the transaction log.

4. (Original) The method of claim 3, wherein said at least one log record characterizes changes made to the primary database in the transaction.

5. (Original) The method of claim 1, wherein said synchronously copying step includes using a file mirroring module.

6. (Original) The method of claim 1, wherein said synchronously copying step includes using file replication hardware.

7. (Original) The method of claim 1, wherein said synchronously copying step includes using file replication software.

8. (Original) The method of claim 1, wherein said synchronously copying step includes synchronously copying information to the transaction log and the mirrored transaction log before completing the transaction at the primary database.

9. (Original) The method of claim 1, wherein said synchronously copying step includes replicating at a file block level the information about the transaction in the transaction log to the mirrored transaction log.

10. (Original) The method of claim 1, further comprising:
copying database schema information from the primary database to a site at which the mirrored transaction log is located to enable transactions to be reconstructed and applied at the replicate database.

11. (Original) The method of claim 10, wherein said generating step includes generating the reconstructed transaction based, at least in part, on said database schema information.

12. (Original) The method of claim 1, wherein said generating step includes formatting the reconstructed transaction so that the reconstructed transaction is in the same format as the transaction at the primary database.

13. (Original) The method of claim 1, wherein said applying step includes verifying that the transaction ordering is correct.

14. (Original) The method of claim 1, wherein said applying step includes applying the reconstructed transaction at the replicate database in the same order as the transaction order at the primary database.

15. (Original) The method of claim 1, further comprising:
responding to a database query at the replicate database while a transaction is being replicated from the primary database to the replicate database.

16. (Original) A computer-readable medium having computer-executable instructions for performing the method of claim 1.

17. (Previously presented) The method of claim 1, further comprising:
downloading a set of computer-executable instructions for performing the method of claim 1.

18. (Previously presented) A system for replicating transactions from a source database to a standby database, the system comprising:

a source database having a transaction log, the transaction log for recording log records for transactions performed at the source database;

a mirrored transaction log for recording mirror copies of the log records for transactions performed at the source database, so as to create at the standby database an exact copy of the transaction log;

a file mirroring module for synchronously replicating log records from the transaction log to the mirrored transaction log as transactions are performed at the source database;

a log reader module for reading log records in the mirrored transaction log and reconstructing transactions for application at the standby database based upon log records in the mirrored transaction log; and

a distribution module for applying the transactions reconstructed by the log reader module at the standby database.

19. (Original) The system of claim 18, wherein said standby database is available for responding to database queries while transactions are being replicated from the source database to the standby database.

20. (Original) The system of claim 18, wherein said transactions include a selected one of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.

21. (Original) The system of claim 18, wherein said log records characterize changes made to the source database based upon transactions performed at the source database.

22. (Original) The system of claim 18, wherein said file mirroring module comprises file replication hardware.

23. (Original) The system of claim 18, wherein said file mirroring module comprises a disk mirroring module.

24. (Original) The system of claim 18, wherein said file mirroring module replicates log records in the transaction log to the mirrored transaction log at a file block level.

25. (Original) The system of claim 18, wherein said file mirroring module replicates log records relating to a particular transaction performed at the source database to the mirrored transaction log before said particular transaction is completed at the source database.

26. (Original) The system of claim 18, wherein said log reader module reconstructs transactions based, at least in part, on database schema information for the source database.

27. (Original) The system of claim 26, further comprising:
database schema information for the source database.

28. (Original) The system of claim 18, wherein said log reader module formats

the reconstructed transactions so that the reconstructed transactions are in the same format as the transaction at the source database.

29. (Original) The system of claim 18, wherein said distribution module applies reconstructed transactions at the standby database in the same order as the order of transactions applied at the source database.

30. (Previously presented) A method for replicating a database operation from a first database to a second database while making the second database available for decision support purposes, the method comprising:

- as a database operation is performed at the first database, generating at least one log record characterizing said operation;

- synchronously recording said at least one log record in a first log associated with the first database and a second log associated with the first log, so that said second log comprises an exact copy of said first log; and

- while the second database is available for decision support purposes, replicating said operation performed at the first database at the second database by performing the substeps of:

- constructing a replicate operation based, at least in part, on said at least one log record in the second log; and

- applying the replicate operation at the second database.

31. (Original) The method of claim 30, wherein said operation includes a selected one of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.

32. (Original) The method of claim 30, wherein said synchronously recording step includes file mirroring.

33. (Original) The method of claim 30, wherein said synchronously recording step includes replicating said at least one log record to the second log at a file block level.

34. (Original) The method of claim 30, wherein said synchronously recording step includes using a disk mirroring module.

35. (Original) The method of claim 30, further comprising:
copying database schema information from the first database prior to performing said operation at the first database.

36. (Original) The method of claim 35, wherein said constructing substep includes constructing a replicate operation based, at least in part, on said database schema information.

37. (Original) The method of claim 35, further comprising:
tracking modifications to said database schema information at the first database;
and
constructing a replicate operation based on said database schema information in effect when the operation is performed at the first database.

38. (Original) The method of claim 30, further comprising:
assigning a unique identifier to database objects at the first database;
if a database object is modified, assigning a different unique identifier to the database object that is modified; and
determining a particular database object to be used in constructing a replicate operation based upon said unique identifier assigned to said particular database object.

39. (Original) The method of claim 30, wherein said constructing substep includes formatting the replicate operation in the same manner as said operation at the first database.

40. (Original) The method of claim 30, wherein said applying substep includes applying the replicate operation at the second database in the same order as said operation

is applied at the first database.

41. (Original) The method of claim 30, wherein making the second database available for decision support purposes includes responding to a database query as said operation is being replicated.

42. (Original) The method of claim 30, wherein making the second database available for decision support purposes includes providing access to data in the second database as said operation is being replicated.

43. (Previously presented) A method for replicating transactions from a primary database to a replicate database while the replicate database remains available for use, the method comprising:

- recording log records for transactions being performed at a primary database in a primary transaction log;

- creating a mirrored transaction log, the mirrored transaction log comprising an exact copy of the log records in the primary transaction log;

- generating reconstructed transactions based on the copies of the log records in the mirrored transaction log; and

- applying the reconstructed transactions at the replicate database while the replicate database remains available for use.

10. EVIDENCE APPENDIX

This Appeal Brief is not accompanied by an evidence submission under §§ 1.130, 1.131, or 1.132.

11. RELATED PROCEEDINGS APPENDIX

Pursuant to Appellant's statement under Section 2, this Appeal Brief is not accompanied by any copies of decisions.